

EXHIBIT 1

(Part 2 of 2)

CISCO'S SUBMISSION OF PROTECTABLE ELEMENTS FROM ITS COPYRIGHTED WORKS

TECHNICAL DOCUMENTATION

Cisco contends that the following excerpts from Cisco's technical documents (also referred to as user manuals, user guides, reference manuals, etc.) are protectable elements of Cisco's copyrighted works. For the avoidance of doubt, where a red box is used to highlight text below, the text within the red box is what Cisco claims to be protectable. As set forth in Cisco's analytic dissection brief, Cisco provides technical documentation that gives users detailed descriptions of the operation of Cisco's user interface. These descriptions were original creations authored by Cisco personnel, and Cisco's personnel could have chosen any arrangement of words and sentences to comprise the text reflected in its technical documents.

show vrrp

To display a brief or detailed status of one or all configured Virtual Router Redundancy Protocol (VRRP) groups on the router, use the **show vrrp** command in privileged EXEC mode.

show vrrp [all | brief]

Cisco IOS IP Application Services Command Reference (2011), at 76.

Usage Guidelines

Use the **ip multicast multipath** command to enable load splitting of IP multicast traffic across multiple equal-cost paths.

If two or more equal-cost paths from a source are available, unicast traffic will be load split across those paths. However, by default, multicast traffic is not load split across multiple equal-cost paths. In general, multicast traffic flows down from the reverse path forwarding (RPF) neighbor. According to the Protocol Independent Multicast (PIM) specifications, this neighbor must have the highest IP address if more than one neighbor has the same metric.

Configuring load splitting with the **ip multicast multipath** command causes the system to load split multicast traffic across multiple equal-cost paths based on source address using the S-hash algorithm. When the **ip multicast multipath** command is configured and multiple equal-cost paths exist, the path in which multicast traffic will travel is selected based on the source IP address. Multicast traffic from different sources will be load split across the different equal-cost paths. Load splitting will not occur across equal-cost paths for multicast traffic from the same source sent to different multicast groups.

Cisco IOS IP Multicast Command Reference (2011), at 293.

Usage Guidelines

Enabling PIM on an interface also enables Internet Group Management Protocol (IGMP) operation on that interface. An interface can be configured to be in dense mode, sparse mode, or sparse-dense mode. The mode describes how the Cisco IOS software populates its multicast routing table and how the software forwards multicast packets it receives from its directly connected LANs. Dense mode interfaces are always added to the table when the multicast routing table is populated. Sparse mode interfaces are added to the table only when periodic join messages are received from downstream routers, or there is a directly connected member on the interface.

Cisco IOS IP Multicast Command Reference (2008), at IMC-233–34

Route Target Extended Community Attribute

The route target (RT) extended community attribute is configured with the **rt** keyword. This attribute is used to identify a set of sites and VRFs that may receive routes that are tagged with the configured route target. Configuring the route target extended attribute with a route allows that route to be placed in the per-site forwarding tables that are used for routing traffic that is received from corresponding sites.

Site of Origin Extended Community Attribute

The site of origin (SOO) extended community attribute is configured with the **soo** keyword. This attribute uniquely identifies the site from which the provider edge (PE) router learned the route. All routes learned from a particular site must be assigned the same site of origin extended community attribute, regardless if a site is connected to a single PE router or multiple PE routers. Configuring this attribute prevents routing loops from occurring when a site is multihomed. The SOO extended community attribute is configured on the interface and is propagated into BGP through redistribution. The SOO should not be configured for stub sites or sites that are not multihomed.

IP Extended Community-List Configuration Mode

Named and numbered extended community lists can be configured in IP Extended community-list configuration mode. To enter IP Extended community-list configuration mode, enter the **ip** **extended-community-list** command with either the **expanded** or **standard** keyword followed by the extended community list name. This configuration mode supports all of the functions that are available in global configuration mode. In addition, you can perform the following operations:

Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IRP-118

Usage Guidelines

Extended community attributes are used to configure, filter, and identify routes for virtual routing and forwarding instances (VRFs) and Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs).

The **match extended-community** command is used to configure match clauses that use extended community attributes in route maps. All of the standard rules of match and set clauses apply to the configuration of extended community attributes.

Cisco IOS IP Routing: EIGRP Command Reference (2011), at 92

Expanded Community Lists

Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to configure patterns to match community attributes. The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in.

Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first. For more information about configuring regular expressions, see the [Regular Expressions](#) appendix of the *Cisco IOS Terminal Services Configuration Guide*.

Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IRP-113–14

Usage Guidelines

The **clear ip bgp** command can be used to initiate a hard reset or soft reconfiguration. A hard reset tears down and rebuilds the specified peering sessions and rebuilds the BGP routing tables. A soft reconfiguration uses stored prefix information to reconfigure and activate BGP routing tables without tearing down existing peering sessions. Soft reconfiguration uses stored update information at the cost of additional memory for storing the updates, to allow you to apply new BGP policy without disrupting the network. Soft reconfiguration can be configured for inbound or outbound sessions.

Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IRP-69

max-metric router-lsa

To configure a router that is running the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the **max-metric router-lsa** command in router configuration mode. To disable the advertisement of a maximum metric, use the **no** form of this command.

max-metric router-lsa [on-startup {seconds | wait-for-bgp}]

no max-metric router-lsa [on-startup {seconds | wait-for-bgp}]

Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IP2R-591

adv-router
[ip-address]

(Optional) Displays all the LSAs of the specified router. If no IP address is included, the information is about the local router itself (in this case, the same as **self-originate**).

link-state-id

(Optional) Portion of the Internet environment that is being described by the advertisement. The value entered depends on the advertisement's LS type. It must be entered in the form of an IP address.

When the link state advertisement is describing a network, the *link-state-id* can take one of two forms:

The network's IP address (as in type 3 summary link advertisements and in autonomous system external link advertisements).

A derived address obtained from the link state ID. (Note that masking a network links advertisement's link state ID with the network's subnet mask yields the network's IP address.)

When the link state advertisement is describing a router, the link state ID is always the described router's OSPF router ID.

When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link state ID is set to Default Destination (0.0.0.0).

Cisco IOS IP Routing Protocols Command Reference, Release 12.4 (2005), at IP2R-613

SNMP notifications can be sent as traps or inform requests. Traps are unreliable because the receiver does not send acknowledgments when it receives traps. The sender cannot determine if the traps were received. However, an SNMP entity that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If the sender never receives the response, the inform request can be sent again. Thus, informs are more likely than traps to reach their intended destination.

Compared to traps, informs consume more resources in the agent and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once; an inform may be tried several times. The retries increase traffic and contribute to a higher overhead on the network.

If you do not enter an **snmp-server host** command, no notifications are sent. To configure the router to send SNMP notifications, you must enter at least one **snmp-server host** command. If you enter the command with no optional keywords, all trap types are enabled for the host.

To enable multiple hosts, you must issue a separate **snmp-server host** command for each host. You can specify multiple notification types in the command for each host.

Cisco IOS IP Switching Command Reference (2011), v. 15.2, at 542

Table 22 *show ip bgp neighbors paths Field Descriptions*

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.

Field	Description
Metric	Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	Autonomous system path for that route, followed by the origin code for that route.

Cisco IOS Multiprotocol Label Switching Command Reference (2011), at 640-41.

Usage Guidelines The **show snmp host** command displays details such as IP address of the Network Management System (NMS), notification type, SNMP version, and the port number of the NMS.
To configure these details, use the **snmp-server host** command.

Command Examples The following is sample output from the **show snmp host** command.

```
Router# show snmp host
Notification host: 10.2.28.6 udp-port: 162 type: inform
user: public security model: v2c
traps: 00001000.00000000.00000000
```

The table below describes the significant fields shown in the display.

Table 5 *show snmp host Field Descriptions*

Field	Description
Notification host	Displays the IP address of the host for which the notification is generated.
udp-port	Displays the port number.
type	Displays the type of notification.
user	Displays the access type of the user for which the notification is generated.
security model	Displays the SNMP version used to send notifications.
traps	Displays details of the notification generated.

Cisco IOS SNMP Support Command Reference (July 2011), at 108–09

dot1x timeout (EtherSwitch)

To set the number of retry seconds between 802.1X authentication exchanges when an Ethernet switch network module is installed in the router, use the **dot1x timeout** command in global configuration mode. To return to the default setting, use the **no** form of this command.

dot1x timeout {**quiet-period** *seconds* | **re-authperiod** *seconds* | **tx-period** *seconds*}

no dot1x timeout {**quiet-period** *seconds* | **re-authperiod** *seconds* | **tx-period** *seconds*}

Syntax Description

quiet-period <i>seconds</i>	Specifies the time in seconds that the Ethernet switch network module remains in the quiet state following a failed authentication exchange with the client. The range is from 0 to 65535 seconds. The default is 60 seconds.
------------------------------------	---

Cisco IOS Security Command Reference, Release 12.4 (2005), at SEC-466.

Usage Guidelines

The security passwords min-length command provides enhanced security access to the router by allowing you to specify a minimum password length, eliminating common passwords that are prevalent on most networks, such as "lab" and "cisco." This command affects user passwords, enable passwords and secrets, and line passwords. After this command is enabled, any password that is less than the specified length will fail.
--

Cisco IOS Security Command Reference, Release 12.4 (2005), at SEC-943.

service sequence-numbers

To enable visible sequence numbering of system logging messages, use the **service sequence-numbers** command in global configuration mode. To disable visible sequence numbering of logging messages, use the **no** form of this command.

service sequence-numbers

no service sequence-numbers

Syntax Description

This command has no arguments or keywords.

Defaults

Disabled.

Command Modes

Global configuration

Command History

Release	Modification
12.0	This command was introduced.

Usage Guidelines

Each system status messages logged in the system logging process have a sequence reference number applied. This command makes that number visible by displaying it with the message. The sequence number is displayed as the first part of the system status message. See the description of the logging commands for information on displaying logging messages.
--

Cisco IOS Configuration Fundamentals Command Reference Release 12.4T (2005), at CF-472.

Usage Guidelines

The command history function provides a record of EXEC commands that you have entered. This function is particularly useful for recalling long or complex commands or entries, including access lists.

To change the number of command lines that the system will record in its history buffer, use the **history size** line configuration command.

The **history** command enables the history function with the last buffer size specified or, if there was not a prior setting, with the default of ten lines. The **no history** command disables the history function.

The **show history EXEC** command will list the commands you have entered, but you can also use your keyboard to display individual commands. Table 34 lists the keys you can use to recall commands from the command history buffer.

Table 34 History Keys

Key(s)	Functions
Ctrl-P or Up Arrow ¹	Recalls commands in the history buffer in a backward sequence, beginning with the most recent command. Repeat the key sequence to recall successively older commands.
Ctrl-N or Down Arrow ¹	Returns to more recent commands in the history buffer after recalling commands with Ctrl-P or the Up Arrow. Repeat the key sequence to recall successively more recent commands.

¹ The arrow keys function only with ANSI-compatible terminals.

Cisco IOS Configuration Fundamentals Command Reference (2010), at CF-237.

Left Arrow ¹ or Ctrl-B	Back character	Moves the cursor one character to the left. When you enter a command that extends beyond a single line, you can press the Left Arrow or Ctrl-B keys repeatedly to scroll back toward the system prompt and verify the beginning of the command entry.
Right Arrow ¹ or Ctrl-F	Forward character	Moves the cursor one character to the right.
Esc, B	Back word	Moves the cursor back one word.
Esc, F	Forward word	Moves the cursor forward one word.
Ctrl-A	Beginning of line	Moves the cursor to the beginning of the line.
Ctrl-E	End of line	Moves the cursor to the end of the command line.

Cisco IOS Configuration Fundamentals Command Reference (2010), at CF-189.

Bridge Assurance is enabled by default and can only be disabled globally. Also, Bridge Assurance can be enabled only on spanning tree network ports that are point-to-point links. Finally, both ends of the link must have Bridge Assurance enabled. If the device on one side of the link has Bridge Assurance enabled and the device on the other side either does not support Bridge Assurance or does not have this feature enabled, the connecting port is blocked.

Cisco NX-OS Layer 2 Switching Configuration Guide (2008), Release 4.0, at 7-3.

A regular expression is entered as part of a command and is a pattern made up of symbols, letters, and numbers that represent an input string for matching (or sometimes not matching). Matching the string to the specified pattern is called pattern matching.

Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at A-1.

\$	Matches the character or null string at the end of an input string.	123\$ matches 0123, but not 1234
*	Matches zero or more sequences of the character preceding the asterisk. Also acts as a wildcard for matching any number of characters.	5* matches any occurrence of the number 5 including none
+	Matches one or more sequences of the character preceding the plus sign.	8+ requires there to be at least one number 8 in the string to be matched
() []	Nest characters for matching. Separate endpoints of a range with a dash (-).	(17)* matches any number of the two-character string 17
	Concatenates constructs. Matches one of the characters or character patterns on either side of the vertical bar.	A(B C)D matches ABD and ACD, but not AD, ABCD, ABBD, or ACCD
-	Replaces a long regular expression list by matching a comma (,), left brace ({), right brace (}), the beginning of the input string, the end of the input string, or a space.	The characters _1300_ can match any of the following strings: <ul style="list-style-type: none"> • ^1300\$ • ^1300space • space1300 • {1300, • ,1300, • {1300} • ,1300,

Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at A-2.

The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first.

Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at A-3.

Syntax	Description
on-startup	(Optional) Configures the router to advertise a maximum metric at startup.
<i>seconds</i>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.
wait-for bgp tag	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.

Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at L3-272.

isis hello-multiplier

To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the **isis hello-multiplier** command in interface configuration mode. To restore the default value, use the **no** form of this command.

isis hello-multiplier multiplier {level-1 | level-2}

no isis hello-multiplier {level-1 | level-2}

Cisco NX-OS Unicast Routing Command Reference (2008), Release 4.0, at L3-224.

IS-IS Overview

IS-IS sends a *hello packet* out every configured interface to discover IS-IS neighbor routers. The hello packet contains information, such as the authentication, area, and supported protocols, which the receiving interface uses to determine compatibility with the originating interface. Compatible interfaces form adjacencies, which update routing information in the link-state database through link-state update messages (LSPs). By default, the router sends a periodic LSP refresh every 10 minutes and the LSPs remain in the link-state database for 20 minutes (the LSP lifetime). If the router does not receive an LSP refresh before the end of the LSP lifetime, the router deletes the LSP from the database.

The LSP interval must be less than the LSP lifetime or the LSPs time out before they are refreshed.

IS-IS Areas

You can design IS-IS networks as a single area that includes all routers in the network or as multiple areas that connect into a backbone or Level 2 area. Routers in a nonbackbone area are Level 1 routers which establish adjacencies within a local area (intra-area routing). Level 2 area routers establish adjacencies to other Level 2 routers and perform routing between Level 1 areas (inter-area routing). A router can have both Level 1 and Level 2 areas configured. These Level 1/Level 2 routers act as area border routers which route information from the local area to the Level 2 backbone area (see Figure 8-1).

Within a Level 1 area, routers know how to reach all other routers in that area. Between areas, routers know how to reach the area border router to get to the Level 2 area. The Level 2 routers know how to reach other area border routers and other Level 2 routers. Level 1/Level 2 routers straddle the boundary between two areas, routing traffic to and from the Level 2 backbone area.

Each IS-IS instance in Cisco NX-OS supports either a single Level 1 or Level 2 area, or one of each. By default, all IS-IS instances automatically support Level 1 and Level 2 routing.

Cisco NX-OS Unicast Routing Configuration Guide, Release 4.0, at 8-2.

PIM Register Messages

PIM register messages are unicast to the RP by designated routers (DRs) that are directly connected to multicast sources. The PIM register message has the following functions:

- To notify the RP that a source is actively sending to a multicast group.
- To deliver multicast packets sent by the source to the RP for delivery down the shared tree.

The DR continues to send PIM register messages to the RP until it receives a Register-Stop message from the RP. The RP sends a Register-Stop message in either of the following cases:

- The RP has no receivers for the multicast group being transmitted.
- The RP has joined the SPT to the source but has not started receiving traffic from the source.

Cisco NX-OS Multicast Routing Configuration Guide (2008), Release 4.0, at 3-7.

If the supplicant is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated supplicant are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the authenticator can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and the supplicant is not granted network access.

Cisco DCNM Security Configuration Guide (2008), Release 4.0, at 6-5.

Changing Global 802.1X Authentication Timers

The following global 802.1X authentication timers are supported on the device:

- Quiet-period timer—When the device cannot authenticate the supplicant, the device remains idle for a set period of time, and then tries again. The quiet-period timer value determines the idle period. An authentication failure might occur because the supplicant provided an invalid password. You can provide a faster response time to the user by entering a number smaller than the default. The default is 60 seconds. The range is from 1 to 65535.

Cisco DCNM Security Configuration Guide (2008), Release 4.0, at 6-14.

Enabling Periodic Reauthentication for an Interface

You can enable periodic 802.1X reauthentication on an interface and specify how often it occurs. If you do not specify a time period before enabling reauthentication, the number of seconds between reauthentication defaults to the global value.

Cisco DCNM Security Configuration Guide (2008), Release 4.0, at 6-14

If the supplicant is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated supplicant are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the authenticator can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and the supplicant is not granted network access.

Cisco NX-OS Security Configuration Guide (2008), Release 4.0, at 7-5.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

Cisco NX-OS System Management Configuration Guide (2008), Release 4.0, at 7-2.

snmp-server enable traps atm pvc

...

Usage Guidelines

SNMP notifications can be sent as traps or inform requests. This command enables both traps and inform requests for the specified notification types. ATM notifications are defined in the CISCO-IETF-ATM2-PVCTRAP-MIB.mib file, available from the Cisco FTP site at <ftp://www.cisco.com/public/mibs/v2/>.

Cisco IOS Asynchronous Transfer Mode Command Reference (2013), at 526.

severity-level

(Optional) The number or name of the desired severity level at which messages should be logged. Messages at or numerically lower than the specified level are logged. Severity levels are as follows (enter the number or the keyword):

[0 | emergencies]—System is unusable
 [1 | alerts]—Immediate action needed
 [2 | critical]—Critical conditions
 [3 | errors]—Error conditions
 [4 | warnings]—Warning conditions
 [5 | notifications]—Normal but significant conditions
 [6 | informational]—Informational messages
 [7 | debugging]—Debugging messages

Cisco IOS Cisco Networking Services Command Reference (2013), at 91.

ip igmp query-interval

Note We recommend that you do not change the default IGMP query interval.

To configure the frequency at which the IGMP querier sends Internet Group Management Protocol (IGMP) host-query messages from an interface, use the **ip igmp query-interval** command in interface configuration mode. To restore the default IGMP query interval, use the **no** form of this command.

```
ip igmp query-interval seconds
no ip igmp query-interval
```

Use the **ip igmp query-interval** command to configure the frequency at which the IGMP querier sends IGMP host-query messages from an interface. The IGMP querier sends query-host messages to discover which multicast groups have members on the attached networks of the router.

Cisco IOS Multicast Command Reference (2013), at 118.

ip msdp mesh-group

To configure a Multicast Source Discovery Protocol (MSDP) peer to be a member of a mesh group, use the **ip msdp mesh-group** command in global configuration mode. To remove an MSDP peer from a mesh group, use the **no** form of this command.

```
ip msdp [vrf vrf-name] mesh-group mesh-name {peer-address| peer-name}
no ip msdp [vrf vrf-name] mesh-group mesh-name {peer-address| peer-name}
```

Cisco IOS Multicast Command Reference (2013), at 225.

A mesh group is a group of MSDP speakers that have fully meshed MSDP connectivity among themselves. Source-Active (SA) messages received from a peer in a mesh group are not forwarded to other peers in the same mesh group.

Cisco IOS Multicast Command Reference (2013), at 226.

Use the **ip multicast multipath** command to enable load splitting of IP multicast traffic across multiple equal-cost paths.

If two or more equal-cost paths from a source are available, unicast traffic will be load split across those paths. However, by default, multicast traffic is not load split across multiple equal-cost paths. In general, multicast traffic flows down from the reverse path forwarding (RPF) neighbor. According to the Protocol Independent Multicast (PIM) specifications, this neighbor must have the highest IP address if more than one neighbor has the same metric.

Configuring load splitting with the **ip multicast multipath** command causes the system to load split multicast traffic across multiple equal-cost paths based on source address using the S-hash algorithm. When the **ip multicast multipath** command is configured and multiple equal-cost paths exist, the path in which multicast traffic will travel is selected based on the source IP address. Multicast traffic from different sources will be load split across the different equal-cost paths. Load splitting will not occur across equal-cost paths for multicast traffic from the same source sent to different multicast groups.

Cisco IOS Multicast Command Reference (2013), at 284.

Enabling PIM on an interface also enables Internet Group Management Protocol (IGMP) operation on that interface. An interface can be configured to be in dense mode, passive mode, sparse mode, or sparse-dense mode. The mode describes how the Cisco IOS software populates its multicast routing table and how the software forwards multicast packets that it receives from its directly connected LANs. Dense mode interfaces are always added to the table when the multicast routing table is populated. Sparse mode interfaces are added to the table only when periodic join messages are received from downstream routers, or there is a directly connected member on the interface.

Cisco IOS Multicast Command Reference (2013), at 330.

ip pim sparse sg-expiry-timer

To adjust the (S, G) expiry timer interval for Protocol Independent Multicast sparse mode (PIM-SM) (S, G) multicast routes (mroutes), use the **ip pim sparse sg-expiry-timer** command in global configuration mode. To restore the default setting with respect to this command, use the **no** form of this command.

```
ip pim [vrf vrf-name] sparse sg-expiry-timer seconds [sg-list access-list]
no ip pim [vrf vrf-name] sparse sg-expiry-timer
```

Cisco IOS Multicast Command Reference (2013), at 405.

Use the **ip pim sparse sg-expiry-timer** command to adjust the expiry timer interval for PIM-SM (S, G) mroute entries to a time value greater than the default expiry timer interval of 180 seconds. This command can be used to lock down the shortest-path tree (SPT) for intermittent sources in PIM-SM network environments, such as sources in trading floor environments that sporadically send financial data streams to multicast groups during trading floor hours.

When a source stops sending traffic to a multicast group, the corresponding (S, G) mroute entry eventually times out and the (S, G) entry is removed. When the source resumes sending traffic to the group, the (S, G) entry is rebuilt. During the short time interval before the (S, G) entry is rebuilt, the traffic is forwarded on the (*, G) forwarding entry. There is a small window of time before the (S, G) entry is completely built in which packets may be dropped. The **ip pim sparse sg-expiry-timer** command can be used to maintain the (S, G) entry so that it will not be removed and the stream will not potentially suffer packet loss.

Cisco IOS Multicast Command Reference (2013), at 406.

Method lists are specific to the type of authorization being requested. AAA supports five different types of authorization:

- **Commands** --Applies to the EXEC mode commands a user issues. Command authorization attempts authorization for all EXEC mode commands, including global configuration commands, associated with a specific privilege level.
- **EXEC** --Applies to the attributes associated with a user EXEC terminal session.

Cisco IOS Security Command Reference: Commands A to C (2013), at 83.

auto	Enables port-based authentication and causes the port to begin in the unauthorized state, allowing only Extensible Authentication Protocol over LAN (EAPOL) frames to be sent and received through the port.
force-authorized	Disables IEEE 802.1X on the interface and causes the port to change to the authorized state without requiring any authentication exchange. The port transmits and receives normal traffic without 802.1X-based authentication of the client. The force-authorized keyword is the default.
force-unauthorized	Denies all access through this interface by forcing the port to change to the unauthorized state, ignoring all attempts by the client to authenticate.

Cisco IOS Security Command Reference: Commands A to C (2013), at 354.

dot1x max-reauth-req

To set the maximum number of times the authenticator sends an Extensible Authentication Protocol (EAP) request/identity frame (assuming that no response is received) to the client, use the **dot1x max-reauth-req** command in interface configuration mode. To set the maximum number of times to the default setting of 2, use the **no** form of this command.

dot1x max-reauth-req *number*

no dot1x max-reauth-req

Cisco IOS Security Command Reference: Commands D to L (2013), at 185.

security passwords min-length

To ensure that all configured passwords are at least a specified length, use the **security passwords min-length** command in global configuration mode. To disable this functionality, use the **no** form of this command.

security passwords min-length *length*

no security passwords min-length *length*

...

The **security passwords min-length** command provides enhanced security access to the device by allowing you to specify a minimum password length, eliminating common passwords that are prevalent on most networks, such as "lab" and "cisco." This command affects user passwords, enable passwords and secrets, and line passwords. After this command is enabled, any password that is less than the specified length will not work.

Cisco IOS Security Command Reference: Commands S to Z at 37 (2013).

show aaa method-lists

To display all the named method lists defined in the authentication, authorization, and accounting (AAA) subsystem, use the **show aaa method-lists** command in user EXEC or privileged EXEC mode.

show aaa method-lists {accounting| all| authentication| authorization}

Syntax Description

accounting	Displays method lists defined for accounting services.
all	Displays method lists defined for all services.
authentication	Displays method lists defined for authentication services.
authorization	Displays method lists defined for authorization services.

Cisco IOS Security Command Reference: Commands S to Z at 185 (2013).

pool *name*

Name of the pool from which global IP addresses are allocated dynamically.

overload

(Optional) Enables the router to use one global address for many local addresses. When overloading is configured, the TCP or User Datagram Protocol (UDP) port number of each inside host distinguishes between the multiple conversations using the same local IP address.

Cisco IOS IP Addressing Services Command Reference (2011), at 440.

ip nat pool

To define a pool of IP addresses for Network Address Translation (NAT), use the **ip nat pool** command in global configuration mode. To remove one or more addresses from the pool, use the **no** form of this command.

ip nat pool *name start-ip end-ip {netmask netmask | prefix-length prefix-length}* [add-route] [type {match-host | rotary}] [accounting *list-name*] [arp-ping] [nopreservation]

no ip nat pool *name start-ip end-ip {netmask netmask | prefix-length prefix-length}* [add-route] [type {match-host | rotary}] [accounting *list-name*] [arp-ping] [nopreservation]

Syntax Description

<i>name</i>	Name of the pool.
<i>start-ip</i>	Starting IP address that defines the range of addresses in the address pool.
<i>end-ip</i>	Ending IP address that defines the range of addresses in the address pool.
<i>netmask netmask</i>	Specifies the network mask that indicates which address bits belong to the network and subnetwork fields and which bits belong to the host field. Specify the netmask of the network to which the pool addresses belong.
<i>prefix-length prefix-length</i>	Specifies the number that indicates how many bits of the netmask are ones (how many bits of the address indicate network). Specify the netmask of the network to which the pool addresses belong.

Cisco IOS IP Addressing Services Command Reference (2011), at 422.

community-string

Password-like community string sent with the notification operation.

Note You can set this string using the **snmp-server host** command by itself, but Cisco recommends that you define the string using the **snmp-server community** command prior to using the **snmp-server host** command.

Note The "at" sign (@) is used for delimiting the context information.

Cisco IOS IP Switching Command Reference (2013), at 526.

max-metric router-lsa

To configure a router that is running the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the **max-metric router-lsa** command in router address family topology or router configuration mode. To disable the advertisement of a maximum metric, use the **no** form of this command.

max-metric router-lsa [**external-lsa** [*max-metric-value*]] [**include-stub**] [**on-startup** {*seconds*|**wait-for-bgp**}] [**summary-lsa** [*max-metric-value*]]

no max-metric router-lsa [**external-lsa** [*max-metric-value*]] [**include-stub**] [**on-startup** {*seconds*|**wait-for-bgp**}] [**summary-lsa** [*max-metric-value*]]

Syntax Description

external-lsa	(Optional) Configures the router to override the external LSA metric with the maximum metric value.
<i>max-metric-value</i>	(Optional) Maximum metric value for LSAs. The configurable range is from 1 to 16777215. The default value is 16711680.
include-stub	(Optional) Configures the router to advertise the maximum metric for stub links in router LSAs.
on-startup	(Optional) Configures the router to advertise a maximum metric at startup.
<i>seconds</i>	(Optional) Maximum metric value for the specified time interval. The configurable range is from 5 to 86400 seconds. There is no default timer value for this configuration option.
wait-for-bgp	(Optional) Configures the router to advertise a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.
summary-lsa	(Optional) Configures the router to override the summary LSA metric with the maximum metric value.

Cisco IOS IP Routing:OSPF Command Reference (2013), at 136.

link-state-id

(Optional) Portion of the Internet environment that is being described by the advertisement. The value entered depends on the advertisement's LS type. It must be entered in the form of an IP address.

When the link state advertisement is describing a network, the *link-state-id* can take one of two forms:

The network's IP address (as in type 3 summary link advertisements and in autonomous system external link advertisements).

A derived address obtained from the link state ID. (Note that masking a network links advertisement's link state ID with the network's subnet mask yields the network's IP address.)

When the link state advertisement is describing a router, the link state ID is always the described router's OSPF router ID.

When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link state ID is set to Default Destination (0.0.0.0).

Cisco IOS IP Routing:OSPF Command Reference (2013), at 185.

Extended community attributes are used to configure, filter, and identify routes for virtual routing and forwarding instances (VRFs) and Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs).

The **match extcommunity** command is used to configure match clauses that use extended community attributes in route maps. All of the standard rules of match and set clauses apply to the configuration of extended community attributes.

Cisco IOS IP Routing: EIGRP Command Reference (2013), at 130.

shutdown (address-family)

To disable the Enhanced Interior Gateway Routing Protocol (EIGRP) address-family protocol for a specific routing instance without removing any existing address-family configuration parameters, use the **shutdown** command in the appropriate configuration mode. To reenab the EIGRP address-family protocol, use the **no** form of this command.

Cisco IOS IP Routing: EIGRP Command Reference (2013), at 276.

Together, a route reflector and its clients form a *cluster*. When a single route reflector is deployed in a cluster, the cluster is identified by the router ID of the route reflector.

The **bgp cluster-id** command is used to assign a cluster ID to a route reflector when the cluster has one or more route reflectors. Multiple route reflectors are deployed in a cluster to increase redundancy and avoid a single point of failure. When multiple route reflectors are configured in a cluster, the same cluster ID is assigned to all route reflectors. This allows all route reflectors in the cluster to recognize updates from peers in the same cluster and reduces the number of updates that need to be stored in BGP routing tables.

Cisco IOS IP Routing: BGP Command Reference (2013), at 74.

bgp router-id

To configure a fixed router ID for the local Border Gateway Protocol (BGP) routing process, use the **bgp router-id** command in router or address family configuration mode. To remove the fixed router ID from the running configuration file and restore the default router ID selection, use the **no** form of this command.

Router Configuration

bgp router-id {*ip-address*| **vrf auto-assign**}

no bgp router-id [**vrf auto-assign**]

Address Family Configuration

bgp router-id {*ip-address*| **auto-assign**}

no bgp router-id

Syntax Description

<i>ip-address</i>	Router identifier in the form of an IP address.
vrf	Configures a router identifier for a Virtual Routing and Forwarding (VRF) instance.
auto-assign	Automatically assigns a router identifier for each VRF.

Command Default

The following behavior determines local router ID selection when this command is not enabled:

- If a loopback interface is configured, the router ID is set to the IP address of the loopback interface. If multiple loopback interfaces are configured, the router ID is set to the IP address of the loopback interface with the highest IP address.
- If no loopback interface is configured, the router ID is set to the highest IP address on a physical interface.

Cisco IOS IP Routing: BGP Command Reference (2013), at 142.

The **clear ip bgp** command can be used to initiate a hard reset or soft reconfiguration. A hard reset tears down and rebuilds the specified peering sessions and rebuilds the BGP routing tables. A soft reconfiguration uses stored prefix information to reconfigure and activate BGP routing tables without tearing down existing peering sessions. Soft reconfiguration uses stored update information, at the cost of additional memory for storing the updates, to allow you to apply new BGP policy without disrupting the network. Soft reconfiguration can be configured for inbound or outbound sessions.

Cisco IOS IP Routing: BGP Command Reference (2013), at 193

distance bgp

To configure the administrative distance for BGP routes, use the **distance bgp** command in address family or router configuration mode. To return to the administrative distance to the default value, use the **no** form of this command.

```
distance bgp external-distance internal-distance local-distance
no distance bgp
```

Syntax Description

<i>external-distance</i>	Administrative distance for external BGP routes. Routes are external when learned from an external autonomous system. The range of values for this argument are from 1 to 255.
<i>internal-distance</i>	Administrative distance for internal BGP routes. Routes are internal when learned from peer in the local autonomous system. The range of values for this argument are from 1 to 255.
<i>local-distance</i>	Administrative distance for local BGP routes. Local routes are those networks listed with a network router configuration command, often as back doors, for the router or for the networks that is being redistributed from another process. The range of values for this argument are from 1 to 255.

Cisco IOS IP Routing: BGP Command Reference (2013), at 271.

Expanded Community Lists

Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to configure patterns to match community attributes. The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first. For more information about configuring regular expressions, see the "Regular Expressions" appendix of the *Terminal Services Configuration Guide*.

Cisco IOS IP Routing: BGP Command Reference (2013), at 324.

Route Target Extended Community Attribute

The route target (RT) extended community attribute is configured with the **rt** keyword. This attribute is used to identify a set of sites and VRFs that may receive routes that are tagged with the configured route target. Configuring the route target extended attribute with a route allows that route to be placed in the per-site forwarding tables that are used for routing traffic that is received from corresponding sites.

Site of Origin Extended Community Attribute

The site of origin (SOO) extended community attribute is configured with the **soo** keyword. This attribute uniquely identifies the site from which the provider edge (PE) router learned the route. All routes learned from a particular site must be assigned the same site of origin extended community attribute, regardless if a site is connected to a single PE router or multiple PE routers. Configuring this attribute prevents routing loops from occurring when a site is multihomed. The SOO extended community attribute is configured on the interface and is propagated into BGP through redistribution. The SOO should not be configured for stub sites or sites that are not multihomed.

Cisco IOS IP Routing: BGP Command Reference (2013), at 330.

Extended community attributes are used to configure, filter, and identify routes for virtual routing and forwarding instances (VRFs) and Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs).

Cisco IOS IP Routing: BGP Command Reference (2013), at 359

Up/Down	The length of time that the BGP session has been in the Established state, or the current status if not in the Established state.
---------	---

Cisco IOS IP Routing: BGP Command Reference (2013), at 821.

State PfxRcd	Current state of the BGP session, and the number of prefixes that have been received from a neighbor or peer group. When the maximum number (as set by the neighbor maximum-prefix command) is reached, the string "PfxRcd" appears in the entry, the neighbor is shut down, and the connection is set to Idle. An (Admin) entry with Idle status indicates that the connection has been shut down using the neighbor shutdown command.
--------------	--

Cisco IOS IP Routing: BGP Command Reference (2013), at 822.

Building the Address Table and Address Table Changes

The device dynamically builds the address table by using the MAC source address of the frames received. When the device receives a frame for a MAC destination address not listed in its address table, it floods the frame to all LAN ports of the same VLAN except the port that received the frame. When the destination station replies, the device adds its relevant MAC source address and port ID to the address table. The device then forwards subsequent frames to a single LAN port without flooding all LAN ports.

Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 10.

- Protocol migration—For backward compatibility with 802.1D devices, 802.1w selectively sends 802.1D configuration BPDUs and TCN BPDUs on a per-port basis.
When a port is initialized, the migrate-delay timer is started (specifies the minimum time during which 802.1w BPDUs are sent), and 802.1w BPDUs are sent. While this timer is active, the device processes all BPDUs received on that port and ignores the protocol type.
If the device receives an 802.1D BPDU after the port migration-delay timer has expired, it assumes that it is connected to an 802.1D device and starts using only 802.1D BPDUs. However, if the 802.1w device is using 802.1D BPDUs on a port and receives an 802.1w BPDU after the timer has expired, it restarts the timer and starts using 802.1w BPDUs on that port.

Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 100

Bridge Assurance

You can use Bridge Assurance to protect against certain problems that can cause bridging loops in the network. Specifically, you use Bridge Assurance to protect against a unidirectional link failure or other software failure and a device that continues to forward data traffic when it is no longer running the spanning tree algorithm.



Note

Bridge Assurance is supported only by Rapid PVST+ and MST.

Bridge Assurance is enabled by default and can only be disabled globally. Also, Bridge Assurance can be enabled only on spanning tree network ports that are point-to-point links. Finally, both ends of the link must have Bridge Assurance enabled. If the device on one side of the link has Bridge Assurance enabled and the device on the other side either does not support Bridge Assurance or does not have this feature enabled, the connecting port is blocked.

Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 175.

**Note**

Do not disable spanning tree on a VLAN unless all switches and bridges in the VLAN have spanning tree disabled. You cannot disable spanning tree on some switches and bridges in a VLAN and leave it enabled on other switches and bridges in the VLAN. This action can have unexpected results because switches and bridges with spanning tree enabled will have incomplete information regarding the physical topology of the network.

Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 108.

The software elects a router as the IGMP querier on a subnet if it has the lowest IP address. As long as a router continues to receive query messages from a router with a lower IP address, it resets a timer that is based on its querier timeout value. If the querier timer of a router expires, it becomes the designated querier. If that router later receives a host query message from a router with a lower IP address, it drops its role as the designated querier and sets its querier timer again.

Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 20 .

PIM Register Messages

PIM register messages are unicast to the RP by designated routers (DRs) that are directly connected to multicast sources. The PIM register message has the following functions:

- To notify the RP that a source is actively sending to a multicast group.
- To deliver multicast packets sent by the source to the RP for delivery down the shared tree.

The DR continues to send PIM register messages to the RP until it receives a Register-Stop message from the RP. The RP sends a Register-Stop message in either of the following cases:

- The RP has no receivers for the multicast group being transmitted.
- The RP has joined the SPT to the source but has not started receiving traffic from the source.

Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 68-69.

Expanded Community Lists

Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to configure patterns to match community attributes. The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first.

Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 274.

max-metric router-lsa (OSPF)

To configure the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the **max-metric router-lsa** command. To disable the advertisement of a maximum metric, use the **no** form of this command.

```
max-metric router-lsa [external-lsa [max-metric-value]] [include-stub] [on-startup [seconds] |
wait-for bgp tag] [summary-lsa [max-metric-value]]
```

```
no max-metric router-lsa [external-lsa [max-metric-value]] [include-stub] [on-startup [seconds] |
wait-for bgp tag] [summary-lsa [max-metric-value]]
```

external-lsa	Specifies the external LSA's.
<i>max-metric-value</i>	(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.
include-stub	Advertises the max-metric for stub links.
on-startup	(Optional) Configures the router to advertise a maximum metric at startup.
<i>seconds</i>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.
wait-for bgp tag	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.
summary-lsa	Specifies the summary LSA's.
<i>max-metric-value</i>	(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.

Defaults Originates router link-state advertisements (LSAs) with normal link metrics.

Command Modes Router configuration
Router VRF configuration

Supported User Roles network-admin
vdc-admin

Command History	Release	Modification
	4.0(1)	This command was introduced.

Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.x (2008), at L3-272.

Use the `ip ospf database` command to display information about different OSPF LSAs.

When the link state advertisement is describing a network, the *link-state-id* argument can take one of two forms:

- The network's IP address (such as Type 3 summary link advertisements and autonomous system external link advertisements).
- A derived address obtained from the link state ID. (Note that masking a network link's advertisement's link state ID with the network's subnet mask yields the network's IP address.)
- When the link state advertisement is describing a router, the link state ID is always the described router's OSPF router ID.
- When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link state ID is set to Default Destination (0.0.0.0).

This command requires the Enterprise Services license.

Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 520.

IS-IS Overview

IS-IS sends a hello packet out every configured interface to discover IS-IS neighbor routers. The hello packet contains information, such as the authentication, area, and supported protocols, which the receiving interface uses to determine compatibility with the originating interface. The hello packets are also padded to ensure that IS-IS establishes adjacencies only with interfaces that have matching maximum transmission unit (MTU) settings. Compatible interfaces form adjacencies, which update routing information in the link-state database through link-state update messages (LSPs). By default, the router sends a periodic LSP refresh every 10 minutes and the LSPs remain in the link-state database for 20 minutes (the LSP lifetime). If the router does not receive an LSP refresh before the end of the LSP lifetime, the router deletes the LSP from the database.

The LSP interval must be less than the LSP lifetime or the LSPs time out before they are refreshed.

IS-IS sends periodic hello packets to adjacent routers. If you configure transient mode for hello packets, these hello packets do not include the excess padding used before IS-IS establishes adjacencies. If the MTU value on adjacent routers changes, IS-IS can detect this change and send padded hello packets for a period of time. IS-IS uses this feature to detect mismatched MTU values on adjacent routers. For more information, see the "Configuring the Transient Mode for Hello Padding" section on page 9-21.

IS-IS Areas

You can design IS-IS networks as a single area that includes all routers in the network or as multiple areas that connect into a backbone or Level 2 area. Routers in a nonbackbone area are Level 1 routers that establish adjacencies within a local area (intra-area routing). Level 2 area routers establish adjacencies to other Level 2 routers and perform routing between Level 1 areas (inter-area routing). A router can have both Level 1 and Level 2 areas configured. These Level 1/Level 2 routers act as area border routers that route information from the local area to the Level 2 backbone area (see Figure 9-1).

Within a Level 1 area, routers know how to reach all other routers in that area. The Level 2 routers know how to reach other area border routers and other Level 2 routers. Level 1/Level 2 routers straddle the boundary between two areas, routing traffic to and from the Level 2 backbone area. Level 1/Level 2 routers use the attached (ATT) bit signal Level 1 routers to set a default route to this Level 1/Level 2 router to connect to the Level 2 area.

In some instances, such as when you have two or more Level 1/Level 2 routers in an area, you may want to control which Level 1/Level 2 router that the Level 1 routers use as the default route to the Level 2 area. You can configure which Level 1/Level 2 router sets the attached bit. For more information, see the "Verifying the IS-IS Configuration" section on page 9-33.

Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 9-2.

NET and System ID

Each IS-IS instance has an associated network entity title (NET). The NET is comprised of the IS-IS system ID, which uniquely identifies this IS-IS instance in the area and the area ID. For example, if the NET is 47.0004.004d.0001.0001.0c11.1111.00, the system ID is 0000.0c11.1111.00 and the area is ID 47.0004.004d.0001.

Designated Intermediate System

IS-IS uses a designated intermediate system (DIS) in broadcast networks to prevent each router from forming unnecessary links with every other router on the broadcast network. IS-IS routers send LSPs to the DIS, which manages all the link-state information for the broadcast network. You can configure the IS-IS priority that IS-IS uses to select the DIS in an area.

**Note**

No DIS is required on a point-to-point network.

Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 9-3.

SNMPv3

SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are as follows:

- **Message integrity**—Ensures that a packet has not been tampered with while it was in-transit.
- **Authentication**—Determines that the message is from a valid source.
- **Encryption**—Scrambles the packet contents to prevent it from being seen by unauthorized sources.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

This section includes the following topics:

- Security Models and Levels for SNMPv1, v2, v3, page 11-4
- User-Based Security Model, page 11-5
- CLI and SNMP User Synchronization, page 11-5

Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 6.x (2013), at 11-3.

Guidelines and Limitations

LLDP has the following configuration guidelines and limitations:

- LLDP must be enabled on the device before you can enable or disable it on any interfaces.
- LLDP is supported only on physical interfaces.
- LLDP can discover up to one device per port.
- LLDP can discover Linux servers, provided they are not using a converged network adapter (CNA). LLDP cannot discover other types of servers.
- DCBXP incompatibility messages might appear when you change the network QoS policy, if a physical loopback connection is in the device. The incompatibility exists for only a short time and then clears.
- DCBXP is not supported for the Cisco Nexus 2000 Series Fabric Extender.
- Beginning with Cisco NX-OS Release 5.2, LLDP is supported for the Cisco Nexus 2000 Series Fabric Extender. LLDP packets can now be sent and received through the Fabric Extender ports for neighbor discovery.
 - All LLDP configuration on Fabric Extender ports occurs on the supervisor. LLDP configuration and show commands are not visible on the Fabric Extender console.
 - LLDP is not supported for a Fabric Extender-virtual port channel (vPC) connection.

Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 6.x (2013), at 18-2.

Configuring Optional LLDP Parameters

You can configure the frequency of LLDP updates, the amount of time for a receiving device to hold the information before discarding it, and the initialization delay time. You can also select the TLVs to include in LLDP packets.

Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 6.x (2013), at 18-7.